



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/603,251	06/25/2003	Thomas J. Heller JR.	POU920020123US1	3683
<div>7590 11/26/2007</div> <div>Lynn L. Augspurger IBM Corporation 2455 South Road, P386 Poughkeepsie, NY 12601</div>				
			EXAMINER PARIKH, KALPIT	
			ART UNIT 2187	PAPER NUMBER
			MAIL DATE 11/26/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 10/603,251	Applicant(s) HELLER ET AL.	
	Examiner Kalpit Parikh	Art Unit 2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 6 September 2007.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-16 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-16 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

I. APPLICATION INFORMATION

The instant application having Application No. 10603251 has a total of 16 claims pending in the application; there is independent claim and 15 dependent claims, all of which are ready for examination by the examiner.

Examiner acknowledges Applicants' amendments to the claims and specification.

Applicants have amended the specification to recite, for related application information, that the invention is related to an application no. 10/603251. The application number for the instant application is 10/603251. Clarification is respectfully requested.

II. REJECTIONS NOT BASED ON PRIOR ART

Claim Rejections - 35 USC § 112

1. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

2. Claims 1-16 rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

The claims are generally narrative and indefinite, failing to conform with current U.S. practice. Claim 1 recites coherent cache regions, cache coherency boundaries, and coherency mode bits, however it is not clear from the claim how these claim elements are interrelated. Claims 6 and 7 each additionally recite logical partitions without clarifying the relationship of logical partitions to the cache coherence regions or the elements of the system. Claim 16 further introduces the term coherency domain and does not clarify how

Art Unit: 2187

coherency domain is related to the previously recited coherent cache regions and coherency boundaries.

Claim 2 recites 'optimizes the cache coherency as supervisory software or firmware expands and contract the number of processors.' The claim does not recite what is necessary to optimize cache coherency as claimed. The cited references recite improving cache coherency operations and therefore the claim limitation is construed as being taught by said references. Clarification is respectfully requested.

Claim 6 recites 'mapping of said logical partitions to allowable physical processors is provided by supervisory firmware of allowable physical processors to an application workload.' It appears the claim intends to recite that the same supervisory software or firmware that provides the mappings of allowable physical processors to an application workload provides mappings of logical partitions to an application workload. Clarification is respectfully requested.

Claim 7 recites 'cache coherence regions.' It is not clear whether the cache coherence region is related to the coherent cache regions recited in claim 1 or a separate element.

Claim 9 recites 'said hardware.' The limitation lacks antecedent basis as the term hardware does not appear prior to the recitation of 'said hardware' in claim 9. While numerous components that may be construed as hardware are previously recited, no specific element is recited as hardware and it is not clear which hardware the limitation is referencing. Clarification is respectfully requested. Clarification is respectfully requested.

Claim 10 recites, "a distinct cache coherency for the address space of the assigned partition as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload in said assigned partition." It

Art Unit: 2187

is not clear how a distinct cache coherency is connected/related to the previously recited claim limitations. Clarification is respectfully requested.

Claim 15 recites 'the buses.' No parent claim recites buses. The limitation lacks antecedent basis.

Claim 16 recites 'a control program for the dispatch of virtual processors for controlling the size and extent of a required coherency domain.' It is not clear whether the claim recites a control program for both the dispatch of virtual processors and for controlling the size and extent of a required domain or if the claim recites virtual processors for controlling the size and extent of a coherency domain.

The foregoing prior art rejections are made in view of the alleged 112 2nd deficiencies in the claims.

III. REJECTIONS BASED ON PRIOR ART

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

4. **CLAIM 1-6,8-11,13-15** rejected under 35 U.S.C. 102(e) as being anticipated by Morioka et al. (US Pat No. 6631447).

As per claim 1, Morioka et al. disclose a multiprocessor computer system (see FIG 1) comprising,

Art Unit: 2187

- a cache coherent symmetric (SMP) computer system of symmetric multiple processors (see e.g., FIG 15: 200: 'CPU') having a plurality of processing nodes (see e.g., FIG 15: 100 'CLUSTER #0') and caches (see e.g., FIG 15: 250 'CACHE') and a node controller (see e.g., FIG 15: 300: 'PMU')
- which use processor state information (see COL 18 LINE 61: 'address translation lookaside buffer in the processor 200') according to mappings provided by supervisor software or firmware of allowable physical processors to an application workload (see COL 14 LINES 'an address translation lookaside map that is controlled by the page allocation function in the operating system')
- to determine which coherent cache regions in the system (see e.g., FIG 6: 'CACHE COHERENCY AREA') are required to examine a coherency transaction produced by a storage request of a single originating processor of said computer system (see COL 18 LINES 59-65: 'Next, it is judged from the information on the cache coherency attribute 223 received from the address translation lookaside buffer in the processor 200 whether or not the target data is of a local cache coherence attribute in step 1302') and
- to change coherency boundaries directly with coherency mode bits (see COL 18 LINE 63: 'When it is of a local attribute,' AND also COL 19 LINE 12: 'On the other hand if an accessed page is of the global cache coherence').

[The LCC/GCC attribute is construed as the coherency mode bits as claimed.

The LCC/GCC attribute determines a coherency boundary for a transaction because the LCC/GCC attribute determines the number of processors for which cache coherency must be guaranteed.]

Art Unit: 2187

As per claim 2, Morioka et al. disclose the multiprocessor computer system according to claim 1, wherein

- a node of said plurality of processing nodes of the computer has dynamic coherency boundaries (see e.g., FIG 7)

[The coherency boundaries change as the operating system assigns new tasks to a node (see PAGE 14: 'controlled by the page allocation function in the operating system').]

- such that hardware of said computer system uses only a subset of the processors in said computer system for a single workload at any specific point in time (see e.g., FIG 7: 'CPU #0' and 'CPU #1')

[A task is construed as a workload as claimed.]

- and optimizes the cache coherency as the supervisor software or firmware expands and contracts the number processors which are being used to run any single workload (see COL 4 LINES 42-46: 'capable of executing cache coherency protocol transactions at a high speed and with minimized interprocessor communications quantities').

As per claim 3, Morioka et al. disclose the multiprocessor computer system according to claim 1,

- wherein multiple instances of a physical node (see e.g., FIG 14: 'CLUSTER #7') are connected with a second level controller (see e.g., FIG 15: PMU in CLUSTER #3) to create a multiprocessor system having multiple node controllers .

[Morioka et al. discloses multiple nodes with multiple node controllers.]

As per claim 4, Morioka et al. disclose the multiprocessor computer system according to claim 1,

Art Unit: 2187

- wherein said node controller uses mode bits (see FIG 5: 315 'LCC/GCC') to determine which processors must receive any given transaction that is received by the node controller (see FIG 5: 321 'AREA LIMITING FUNCTION').

As per claim 5, Morioka et al. disclose the multiprocessor computer system according to claim 1

- wherein a second level controller (see e.g., FIG 14: 300 'PMU' in CLUSTER #7) is provided which uses mode bits (see FIG 14: 310) to determine which nodes must receive any given transaction that is received by the second level controller.

As per claim 6, Morioka et al. disclose the multiprocessor computer system according to claim 1,

- wherein logical partitions are provided (see FIG 2: 'PHYSICAL ADDRESS SPACE') and mapping of said logical partitions to allowable physical processors is provided (see FIG 2: 'LOCAL SHARED MEMORY #0 (CLUSTER #0)') by supervisor software or firmware of allowable physical processors to an application workload.

[The address space is partitioned among the clusters and mapped to individual clusters.]

As per claim 8, Morioka et al. disclose the multiprocessor computer system according to claim 1

- wherein a single workload (see FIG 7: 'TASK #0') uses only a subset of the total processors in the computer system (see e.g., FIG 7: CPU #0, CPU #1) for a single workload (see FIG 7: 'TASK #0') at any specific point in time for an assigned partition (see e.g., FIG 7: CLUSTER #0) and a distinct cache coherency for the address space (see FIG 7: 'LCC') of the assigned partition as the supervisor software or firmware

Art Unit: 2187

expands and contracts the number of processors which are being used to run any single workload in said assigned partition (see COL 14 LINES 57-63: 'page allocation function in the operating system').

[Local cache coherency is a distinct cache coherency for the address space of the assigned partition.]

As per claim 9, the multiprocessor computer system according to claim 1

- wherein a single workload uses only a subset of the total processors in the computer system for a single workload at any specific point in time (see e.g., FIG 7: TASK #0), and multiple cache coherent regions are assigned for different partitions (see e.g., FIG 7: 'CLUSTER #0') as more independent workloads (see e.g., FIG 7: 'TASK #1') coexist on said hardware.

[Each local cache coherent region of a cluster is considered to be a different cache coherent region. Morioka et al. discloses independent workloads can be assigned to the hardware (see e.g., FIG 7: TASK #0 and TASK #1).]

As per claim 10, the multiprocessor computer system according to claim 1 wherein

- cache coherence regions encompass subsets of processors and caches in the computer system (see e.g., FIG 14: LCC CLUSTER #0) and a single workload (see FIG 7: 'TASK #0') uses only a subset of the total processors (see e.g., FIG 7: 'CPU #0') in the computer system for a single workload (see e.g., FIG 7: TASK #0) at any specific point in time for an assigned partition (see e.g., FIG 7: 'CLUSTER # 0') and
- a distinct cache coherency for the address space of the assigned partition (see FIG 7: 'LCC') as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload in said assigned partition).

Art Unit: 2187

[Local cache coherency is a distinct cache coherency for the address space of the assigned partition.]

As per claim 11, Morioka et al. disclose the multiprocessor computer system according to claim 1

- wherein software and/or firmware (see COL 14 LINES 57-63: 'operating system') define which subset of processors in multiprocessor must participate in a coherency transaction (see COL 14 LINES 57-63: 'page allocation function in the operating system') independent of which processing node is connected to the physical DRAM storage being requested by said single originating processor

[The determination to decide which cache coherency regions must participate in cache coherency is determined based on the LCC/GCC attribute (i.e., independent of which processing node is connected to the physical DRAM).]

As per claim 13, Morioka et al. disclose the multiprocessor computer system according to claim 1

- wherein cache coherence mode bits are appended to a processors storage transactions when transmitted to a connected processor of said multiprocessor computer system (see e.g., FIG 5 "LCC/GCC-REAL ADDRESS").

As per claim 14, Morioka et al. disclose the multiprocessor computer system according to claim 13

- wherein said cache coherence mode bits are used in a decision determining whether the single originating processor's storage request must be transmitted to additional processors in the system (see e.g., FIG 11: 1201).

Art Unit: 2187

As per claim 15, Morioka et al. disclose the multiprocessor computer system according to claim 14

- wherein an increase in the effective utilization of the address bandwidth of the buses used to interconnect the processors of a multiprocessor system allows movement of workload among physical processors in a multiprocessor system at the same time as a reduction of the address bandwidth required to maintain cache coherency among all the processors is caused.

[The claim appears to be reciting the benefit of the invention (see Specification PAGE 8 LINES 3-10). Such benefits are understood to be present in the system of Morioka et al. (see Morioka et al. (ABSTRACT)).]

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. **CLAIMS 12** rejected under 35 U.S.C. 103(a) as being unpatentable over Morioka et al. (US Pat No. 6631447) in view of Hagersten et al. (US Pat No. 6226671).

As per claim 12, Morioka et al. disclose the multiprocessor computer system according to claim 11.

However, Morioka et al. do not expressly disclose wherein

- movement of a process between nodes of said symmetric multiple processors of said multiprocessor is effectuated without moving physical storage contents and without requiring subsequent broadcasting of storage references originated by the process

Art Unit: 2187

from said single originating processor's storage request to all of the caches in the multiprocessor.

In the same field of endeavor Hagersten et al. disclose

- movement of a process between nodes of said symmetric multiple processors of said multiprocessor (see Hagersten et al. COL 5 LINES 62-65: 'a process migrates from one node in the multiprocessor architecture to another node') is effectuated without moving physical storage contents (see Hagersten et al. COL 6 LINES 17-20: 'An alternative to copying the data from one to another is to change the translation address from local to global') and
- without requiring subsequent broadcasting of storage references originated by the process from said single originating processor's storage request to all of the caches in the multiprocessor (see Hagersten et al. COL 7 LINES 37-39: 'a method of dynamically changing the designation of memory blocks is desirable').

[Hagersten et al. teaches that while initially the storage references are treated global after a process has migrated to a remote node, subsequent processing may cause the storage reference to be treated local. Subsequent storage references would not require global broadcast.]

Hagersten et al. and Morioka et al. are analogous art because they are from the same field of endeavor namely memory access and control.

It would have been obvious to a person of ordinary skill in the art to modify Morioka et al. to perform process migrations without moving physical storage contents.

The suggestion/motivation for doing so would have been because Hagersten et al. disclose it is desirable to avoid process migrations that require movement of data blocks (see Hagersten et al. COL 6 LINES 16-17).

Therefore it would have been obvious to a person of ordinary skill in the art to modify Morioka et al. to perform process migration across nodes without moving physical storage contents as taught by Hagersten et al. (see Hagersten et al. COL 6 LINES 21-30) for the benefit of process migration without movement of data blocks to arrive at the invention as specified in the claims.

IV. ALLOWABLE SUBJECT MATTER

Claim 7 and 16 indicated allowable over prior art of record.

Claims 7 and 16 would be allowable over the prior art of record if rewritten to overcome the rejection(s) under 35 U.S.C. 112, 2nd paragraph, set forth in this Office action and to include all of the limitations of the base claim and any intervening claims.

The following is a statement of reasons for the indication of allowable subject matter:

The primary reasons for allowance of claims 7 in the instant application is the combination with the inclusion in these claims of

- a hypervisor assigns cache coherence regions which encompass subsets of said processors and caches in the system chosen for their physical proximity and defines a distinct cache coherency region for each of said logical partitions.

The prior art of record neither anticipates nor renders obvious the above-recited combination.

The primary reasons for allowance of claims 16 in the instant application is the combination with the inclusion in these claims of

Art Unit: 2187

- a control program for the dispatch of virtual processors for controlling the size and extent of a required coherency domain changes said coherency boundaries directly with coherency mode bits.

The prior art of record neither anticipates nor renders obvious the above-recited combination.

IMPORTANT NOTE:

If the applicant should choose to rewrite the independent claims to include the limitations recited in claims [7 and 16], the applicant is encouraged to amend the **title of the invention** such that it is descriptive of the invention as claimed as required by sec. 606.01 of the MPEP. Furthermore, the **Summary of the Invention** and the **Abstract** should be amended to bring them into harmony with the allowed claims as required by paragraph 2 of sec. 1302.01 of the MPEP.

As allowable subject matter has been indicated, applicant's response must either comply with all formal requirements or specifically traverse each requirement not complied with. See 37

C.F.R. § 1.111(b) and § 707.07(a) of the M.P.E.P.

V. ACKNOWLEDGMENT OF ISSUES RAISED BY THE APPLICANT

Applicants' arguments filed September 6, 2007 have been fully considered but they are not deemed persuasive with regard to Morioka et al. (US Pat No. 6631447). A response to Applicants' argument appears below.

Applicant's arguments, see Remarks PAGE 8, filed September 6, 2007, with respect to rejections under Hagesrten (6226671) have been fully considered and are persuasive. The rejection of claims 1,6 and 11-12 as being anticipated by Hagersten has been withdrawn. The rejection of claim 7 over Hagersten and IBM (NPL) has been withdrawn.

RESPONSE TO AMENDMENTS/ARGUMENTS

1st POINT OF ARGUMENT:

Applicants argue 'Morioka did not teach the use of virtual processor dispatch (see new claim 16) and focused only on manipulation of page table information that is cached in the TLB.'

Examiner notes Morioka et al. were not relied upon to disclose such matter.

2nd POINT OF ARGUMENT:

Applicants argue "This invention is much more direct, the control program can change the coherency boundaries directly by changing our coherency mode bits which is claimed in claim 1. The claimed invention does not require the intermediate steps of changing page tables and manipulating TLB entries."

The apparent differences between the prior art and the invention, as pointed out by Applicants' representatives, are differences in the manner in which the coherency mode bits (i.e., LCC/GCC of Morioka et al.) are changed to change the coherency boundaries. The apparent differences are not reflected in the claim 1 because the claim 1 does not recite how the coherency mode bits are changed.

Morioka et al. discloses the coherency boundary is directly dependent on the LCC/GCC attributes (see e.g., FIG 8: 902).

VI. CLOSING COMMENTS

RELEVANT ART CITED BY THE EXAMINER

The following prior art made of record and not relied upon is cited to establish the level of skill in the applicant's art and those arts considered reasonably pertinent to applicant's disclosure. See MPEP 707.05(c).

NPL: The Written Opinion of the International Searching Authority for PCT 2004/050878 which claims benefit to the instant application.

STATUS OF CLAIMS IN THE APPLICATION

The following is a summary of the treatment and status of all claims in the application as recommended by M.P.E.P. ' 707.07(i):

Art Unit: 2187

Via. CLAIMS REJECTED IN THE APPLICATION

Per the instant office action, claims 1-16 have received a second action on the merits and are subject of a non-final office action.

For at least the above reasons it is the examiner's position that the applicant's claims are not in condition for allowance.

Art Unit: 2187

VII. DIRECTION OF FUTURE CORRESPONDENCES

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Kalpit Parikh whose telephone number is (571) 270-1173. The examiner can normally be reached on MON THROUGH FRI 7:30 TO 5:00.

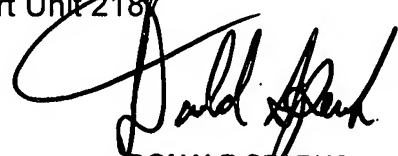
If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on (571) 272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Kalpit Parikh/

November 16, 2007

Kalpit Parikh
Examiner
Art Unit 2187



DONALD SPARKS
SUPERVISORY PATENT EXAMINER